

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: WILLIAMS JR. ET AL.

Application No. 10/811,044

Confirmation No. 9536

Filed: March 27, 2004

Group Art Unit: 2165

Examiner: HICKS, Michael J.

For: Bypassing Native Storage Operations
By Communicating Protected Data
Within Locking Messages Using a
Lock Manager Independent of the
Storage Mechanism

APPEAL BRIEF

Mail Stop Appeal Brief-Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

On November 19, 2007, Appellants appealed from the final Office action mailed July 23, 2007. Appellants submit this Appeal Brief with the fee for submitting this Appeal Brief being provided via EFS-Web. Appellants have also submitted a petition for an extension of time. Appellants further authorizes the Commissioner to charge Deposit Account No. 501430 with any additional fees due in connection with the submission of this paper, and petitions for any additional extension of time which may be deemed necessary. Moreover, Appellants submit that the prior art of record neither teaches nor suggests all the limitations of any claim, and therefore, Appellants request all rejections be reversed and all claims be allowed.

(i) REAL PARTY IN INTEREST

The above-identified application has been assigned to Cisco Technology, Inc. by all inventors, with this assignment recorded in the USPTO at Reel 015159, Frame 0873, with a recordation date of March 27, 2004.

(ii) RELATED APPEALS AND INTERFERENCES

None.

(iii) STATUS OF CLAIMS

Claims 1-26 are pending in the application.

No claims stand as canceled.

No claims stand as objected to.

Claims 1-26 stand as rejected.

Claims 1-26 are on appeal in the application.

(iv) STATUS OF AMENDMENT

NONE.

(v) SUMMARY OF CLAIMED SUBJECT MATTER

There are seven independent claims on appeal. Independent claims 1, 6, 8, 10, 12, 17 and 22, in various claim formats and reciting various limitations surrounding a lock manager and its use in communicating protected data within locking messages. There are many embodiments described in the extensive specification and illustrated in the large number of figures, and only one or some of these embodiments is/are described herein in relation to each independent claim on appeal as required by the Rules.

Independent claim 1 is an apparatus claim for protecting data using locks reciting patentably distinct limitations to which Appellants especially refer the Board to the apparatus illustrated in FIG. 5A, which is described at least on page 19, lines 17-25, and the message sequence chart of FIG. 2, which is described at least from page 14, line 23 to page 18, line 3. Claim 1 recites a lock manager (501 of FIG. 5, 200 of FIG. 2), a plurality of requesters (511-519 of FIG. 5, 204-208 of FIG. 2). The lock manager is configured to receive lock requests from each of the plurality of requesters (211, 212, 213 of FIG. 2), and to selectively grant said lock requests (221, 231, 241 of FIG. 2) which includes communicating grants from the lock manager to the plurality of requesters, wherein at least one of said communicated grants includes said protected data (231, 241 of FIG. 2).

Independent claim 6 is a method claim for protecting data using locks reciting patentably distinct limitations to which Appellants especially refer the Board to the apparatus illustrated in FIG. 5A, which is described at least on page 19, lines 17-25, the message sequence chart of FIG. 2, which is described at least from page 14, line 23 to page 18, line 3, and the flow diagrams of FIGs. 4A-B, which are described at least from page 18, lines 4-27. Claim 1 recites operations performed by a lock manager (501 of FIG. 5, 200 of FIG. 2) controlling access to protected data maintained in native storage independent of the lock manager, wherein the lock manager does not access said protected data from said native storage. The method comprises: receiving a release of a lock for use in controlling access to said protected data, the release including said protected data (225, 233 of FIG. 2, 422, 436 of FIG. 4B); identifying a next requester to be granted the lock in response to said receiving the release of the lock (428 of FIG. 4B); copying said protected data from the release into a grant message (438 of FIG. 4B); and sending the grant message to the next requester, the grant message including said protected data (231, 241 of FIG. 2, 442 of FIG. 4B).

Independent claim 8 is a Beauregard-style claim (directly corresponding to method claim 6) reciting steps for protecting data using locks reciting patentably distinct limitations to which Appellants additionally refer the Board to FIG. 5B, described at least from page 19, line 25 to page 20, line 20. Claim 8 recites a computer-readable medium (542, 543 of FIG. 5B) tangibly storing thereon computer-executable instructions for performing steps by a lock manager for controlling access to protected data maintained in native storage independent of the lock manager, wherein the lock manager does not access said protected data from said native storage. The recited steps are described *supra* in relation to independent claim 6.

Independent claim 10 is an apparatus claim reciting means plus function language directly based on the method operations of independent claim 6. Appellants especially refer the Board to the apparatus illustrated in FIG. 5A, which is described at least on page 19, lines 17-25, the apparatus illustrated in FIG. 5B, described at least from page 19, line 25 to page 20, line 20, the message sequence chart of FIG. 2, which is described at least from page 14, line 23 to page 18, line 3, and the flow diagrams of FIGs. 4A-B, which are described at least from page 18, lines 4-27. Claim 10 recites: A lock manager (501 of FIG. 5, 540 of FIG. 5B, 200 of FIG. 2) controlling access to protected data maintained in native storage independent of the lock manager, wherein the lock manager does not access said protected data from said native storage. The lock manager comprises: means for receiving a release of a lock for use in controlling access to said protected data, the release including said protected data (225, 233 of FIG. 2, 422, 436 of FIG. 4B); means for identifying a next requester to be granted the lock in response to said receiving the release of the lock (428 of FIG. 4B); and means for copying said protected data from the release into a grant message and for sending the grant message to the next requester (231, 241 of FIG. 2, 438, 442 of FIG. 4B).

Independent claim 12 is a method claim for protecting data using locks reciting patentably distinct limitations to which Appellants especially refer the Board to the apparatus illustrated in FIG. 5A, which is described at least on page 19, lines 17-25, the message sequence chart of FIG. 2, which is described at least from page 14, line 23 to page 18, line 3, and the flow diagrams of FIGs. 4A-B, which are described at least from page 18, lines 4-27. Claim 12 recites operations performed by a lock manager (501 of FIG. 5, 200 of FIG. 2) controlling access to protected data maintained in native storage independent of the lock manager, wherein the lock manager does not access said protected data from said native storage. The method comprises: receiving locking requests (211, 212 of FIG. 2, 402 of FIG. 4A) for a lock controlling access to said protected data from a first requester (204 FIG. 2) and a second requester (206 of FIG. 2); sending a first grant message to the first requester, the first grant message not including said protected data, and in response to identifying one or more requesters is waiting for the lock after the first requester, including an indication to return said protected data in the first grant message (221 of FIG. 2, 430, 432, 436, 440, 442 of FIG. 4B); and receiving a first release message corresponding to the first grant message for the lock from the first requester, the first release message including said protected data (225 of FIG. 2, 422, 436 of FIG. 4B).

Independent claim 17 is a Beauregard-style claim (directly corresponding to method claim 12) reciting steps for protecting data using locks reciting patentably distinct limitations to which Appellants additionally refer the Board to FIG. 5B, described at least from page 19, line 25 to page 20, line 20. Claim 17 recites a computer-readable medium (542, 543 of FIG. 5B) tangibly storing thereon computer-executable instructions for performing steps by a lock manager for controlling access to protected data maintained in native storage independent of the lock manager, wherein the lock manager does not access said protected data from said native storage. The recited steps are described *supra* in relation to independent claim 12.

Independent claim 22 is an apparatus claim reciting means plus function language directly based on the method operations of independent claim 12. Appellants especially refer the Board to the apparatus illustrated in FIG. 5A, which is described at least on page 19, lines 17-25, the apparatus illustrated in FIG. 5B, described at least from page 19, line 25 to page 20, line 20, the message sequence chart of FIG. 2, which is described at least from page 14, line 23 to page 18, line 3, and the flow diagrams of FIGs. 4A-B, which are described at least from page 18, lines 4-27. Claim 10 recites: A lock manager (501 of FIG. 5, 540 of FIG. 5B, 200 of FIG. 2) lock manager controlling access to protected data maintained in native storage independent of the lock manager, wherein the lock manager does not access said protected data from said native storage. The lock manager comprises: means for receiving locking requests (211, 212 of FIG. 2, 402 of FIG. 4A) for a lock controlling access to said protected data from a first requester (204 FIG. 2) and a second requester (206 of FIG. 2); means for sending a first grant message to the first requester, the first grant message not including said protected data, and in response to identifying one or more requesters is waiting for the lock after the first requester, including an indication to return said protected data in the first grant message (221 of FIG. 2, 430, 432, 436, 440, 442 of FIG. 4B); and means for receiving a first release message for the lock from the first requester, the first release message including said protected data (225 of FIG. 2, 422, 436 of FIG. 4B).

(vi) GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

The issue presented on appeal is listed below, and then addressed in corresponding headings and subheadings hereinafter. Although there are additionally reasons that all claims are patentably distinct over the prior art of record, Appellants have elected solely for the purposes of this Appeal Brief to limit the issues to the issue listed below and discussed *infra*. Appellants respectfully request the Board reverse all rejections.

Whether all pending claims, claims 1-26, are unpatentable for the stated rejections in the final Office mailed July 23, 2007, as being obvious under 35 USC § 103(a) over the Amiri et al., "Highly Concurrent Shared Storage" article in view of the Yun et al., "An Efficient Lock Protocol for Home-based Lazy Release Consistency" article.

(vii) ARGUMENT

Appellants present the arguments herein based on the following outline. Although there are additionally reasons that all claims are patentably distinct over the prior art of record, Appellants have elected solely for the purposes of this Appeal Brief to limit the issues to the issue listed below and discussed *infra*. Appellants respectfully request the Board reverse all rejections.

Claims 1-26 are patentably distinct over prior art of record, and Appellants request the Board reverse the rejections of all claims under 35 USC § 103(a) as being obvious over the Amiri et al., "Highly Concurrent Shared Storage" article in view of the Yun et al., "An Efficient Lock Protocol for Home-based Lazy Release Consistency" article.

- (a) The Office fails to present a *prima facie* rejection of any of claims 1-5.
- (b) The Office fails to present a *prima facie* rejection of any of claims 6-11.
- (c) The Office fails to present a *prima facie* rejection of any of claims 12-26.
- (d) The Office fails to present a *prima facie* rejection of any of claims 4, 8, 9, 11, 13-16, 18-21, and 23-26

(vii)(a) The Office fails to present a *prima facie* rejection of any of claims 1-5, as the prior art of record, alone or in combination, neither teaches nor suggests a grant message from a lock manager to a requester including the protected data.

Group: independent claim 1 and its dependent claims 2-5.

Representative claim: independent claim 1.

The disclosure in the originally filed application describes protecting data using locks with the ability to pass the protect data through the lock manager, using the constructs and mechanisms described therein. Locks are commonly used to provide distributed applications with a means to synchronize their accesses to shared resources, such as described in relation to prior art system of FIG. 1. Note, the prior art lock manager receives lock requests and lock releases, and grants requests to access the protected data. The lock manager does not access, receive, nor process the protected data.

All claims recite patentably distinct limitations over the prior art of record, in that the lock manager, which does not have access to the native storage of the protected data (i.e., it is actually a lock manager, not a process processing the native data), provides a conduit to communicate the protected data.

The following three paragraphs, extracted from page 13, line 22 to page 14, line 22 of the original disclosure, describe the operation of one embodiment.

"One embodiment provides an indirect interprocess communication bypass channel through a lock mechanism that is connected to many processors. These processors normally communicate through shared global memory and use locks to enforce coherency. Under certain conditions, data can be transferred through lock messages instead of going through shared global memory. The data can be piggy-backed to the lock release message, go through the lock mechanism, and be piggy-backed to the lock grant message. The data is not stored for any significant amount of time in the lock mechanism. The lock messages typically include control signals to indicate when the conditions are right to use the bypass channel. This enforces the coherency of the shared memory location that might be bypassed.

When claiming to piggy-back the data to the lock message, the bypass channel could be either serial or parallel to the lock message channel, as long as there is a strong binding of lock messages to bypass data. In one embodiment, when requesting a lock, the request message includes an indication if it is willing to accept data through the bypass channel. When the lock is finally granted, the grant message indicates if it has data in the bypass channel, and if there is an entry following it in the locking queue that is willing to accept data through the bypass channel. If the grant indicates that data is present in the bypass channel, then the critical section can skip the read of the global shared memory location and use the data from the bypass channel instead.

If the grant indicates that the next entry in the locking queue is willing to accept data from the bypass channel, then the critical section of code can skip the write of the global shared memory location and send the data through the bypass channel instead. The critical section of code can always send the data through the bypass channel with the hope that a new arrival in the locking queue can use the data, but it must first commit the write to global shared memory if it is not certain. When the lock is released, an indication is made in the release message whether the bypass channel has data in it or not. The data in the bypass channel is typically not stored in memory in the lock mechanism; rather it is simply copied (possibly using a temporary storage location or register) from the release message and attached to the subsequent grant message."

Turning to claims, independent claim 1 recites:

"1. An apparatus for protecting data using locks, the apparatus comprising:
a lock manager configured to control access via a lock to protected data maintained in native storage independent of the lock manager, wherein the lock manager does not access said protected data from said native storage; and
a plurality of requesters;
wherein the lock manager is configured to receive lock requests for the lock from each of the plurality of requesters, and to selectively grant said lock requests which includes communicating grants from the lock manager to the plurality of requesters, wherein at least one of said communicated grants includes said protected data."

Claims 1-5 stand rejected per the final Office mailed July 23, 2007, as being obvious under 35 USC § 103(a) over the Amiri et al., "Highly Concurrent Shared Storage" article in view of the Yun et al., "An Efficient Lock Protocol for Home-based Lazy Release Consistency" article. Appellants respectfully traverse the rejection of these claims, as the prior art of record neither teaches nor suggests all recited limitations of claim 1.

It is well-established law that the burden is on the Office to initially present a *prima facie* unpatentability rejection, before Applicant has any burden of proof of disproving any application of a cited reference against a claim. *In re Warner*, 379 F2d. 1011, 1016, 154 USPQ 173, 177 (C.C.P.A. 1967); *Ex parte Skinner*, 2 USPQ2d 1788, 1788-89 (B.P.A.I. 1986). The reference(s) *must teach each and every aspect of the claimed invention* either explicitly or impliedly, and the burden is on the Office to present a *prima facie* case of unpatentability.

The Office relies on Amiri et al. for teaching most limitations, and admits that "Amiri fails to disclose at least one of said communication grants includes said protected data." Final Office action, page 7, lines 1-2. Appellants agree. The Office then relies on Yun et al. for its teaching that a releasing process sends diffs with write notices as a lock grant message.

However, the limitation to which the Office admits that Amiri et al. fails to teach is *the lock manager communicating grants including said protected data from the lock manager to the requesters*. A releasing process of Yun et al. is not a lock manager. Moreover, the releasing process of Yun et al. has access to the protective data in its native storage (it reads, writes, and processes the data), which conflicts with the recited limitation of "wherein the lock manager does not access said protected data from said native storage."

Yun et al. is an article describing an enhancement to home-based lazy release consistency (HLRC). It does not describe fully HDRC, but rather assumes the reader understands the basics of the system. A lock manager is used by the system described by Yun et al., with this lock manager providing the synchronization among processes and allows a process to acquire the data protected by the lock, with the lock manager being separate and distinct from an process P0, P1, P2, P3. Yun et al. refers to an "acquirer" or "acquiring process"; and a process becomes such based on the determination of the lock manager.

Appellants refer the Board to page 527 of Yun et al., which states that it is an implementation of HDRC described reference [13] (Y. Zhou et al, "Performance Evaluation of Two Home-Based Lazy Release Consistency Protocols for Shared Virtual Memory Systems"). The first paragraph of section 3.5 of Y. Zhou et al. clarifies the operation of Yun et al, with this paragraph being reproduced below.

"Synchronization handling and related coherence checking for all four prototypes is implemented at user level using NX/2 messages. Each lock has a manager, which is assigned in a round-robin fashion among the processors. The manager keeps track of which processor has most recently requested the lock. All lock acquire requests are sent to the manager unless the node itself holds the lock. The manager forwards the lock request to the processor that last requested the lock. The request message contains the current maximal vector timestamp of the acquiring processor. When the lock arrives, it contains the

releaser's knowledge of all time intervals for the requester to update its timestamp vectors."

Therefore, the system of Yun et al. includes process P0, P1, P2, P3 as well as a "lock manager" for the protected data. In fact, Y. Zhou et al. calls it a "manager" of the lock. The use of the term lock manager in the present disclosure is consistent with the usage of manager of locks by Y. Zhou et al. and hence with Yun et al.; and there is no teaching nor a suggestion in Yun et al. that the lock manager communicates the protected data. In fact, there is no teaching nor a suggestion in Yun et al. that its lock manager actually ever receives the protected data.

Rather, Yun et al.'s lock manager communicates a grant message *without the protected data* to the process currently holding the lock. When the current process releases the lock, it forwards the diffs to newly acquiring process. However, this neither teaches nor suggests the limitation that *the lock manager communicates grants including said protected data from the lock manager to the requesters*. A releasing process of Yun et al. is not a lock manager. Moreover, the releasing process of Yun et al. has access to the protective data in its native storage (it reads, writes, and processes the data), so it being equated to the recited limitation of "a lock manager" conflicts with the recited limitation of "wherein the lock manager does not access said protected data from said native storage."

Appellants further traverse the combination of Yun et al. with Amiri et al. that would render a claim unpatentable as being obvious as Yun et al. communicating of data is from a releasing process to an acquiring process; and Yun et al. or Amiri et al., alone or in combination, neither teach nor suggest protected data being communicated or even touched by a lock manager as recited in every pending claim.

For at least these reasons, the Office has failed to present a *prima facie* rejection of independent claim 1, and Appellants request the Board reverse the rejections of independent claim 1 and dependent claims 2-5.

(vii)(b) The Office fails to present a *prima facie* rejection of any of claims 6-11, as the prior art of record, alone or in combination, neither teaches nor suggests a lock manager receiving a release including the protected data, copying the protected data from the release into a grant message, and sending a grant message with the copied protected data to the identified next requester.

Group: claims 6-11. Representative claim: independent claim 6.

Independent claim 6 recites:

"6. A method performed by a lock manager controlling access to protected data maintained in native storage independent of the lock manager, wherein the lock manager does not access said protected data from said native storage, the method comprising:

receiving a release of a lock for use in controlling access to said protected data, the release including said protected data;

identifying a next requester to be granted the lock in response to said receiving the release of the lock;

copying said protected data from the release into a grant message; and

sending the grant message to the next requester, the grant message including said protected data."

Claims 6-11 stand rejected per the final Office mailed July 23, 2007, as being obvious under 35 USC § 103(a) over the Amiri et al., "Highly Concurrent Shared Storage" article in view of the Yun et al., "An Efficient Lock Protocol for Home-based Lazy Release Consistency" article. Appellants respectfully traverse the rejection of these claims, as the prior art of record neither teaches nor suggests all recited limitations of any of these claims. Note, independent claims 8 and 10 are different claim formats directly based on claim 6.

It is well-established law that the burden is on the Office to initially present a *prima facie* unpatentability rejection, before Applicant has any burden of proof of disproving any application of a cited reference against a claim. *In re Warner*, 379 F2d. 1011, 1016, 154 USPA 173, 177 (C.C.P.A. 1967); *Ex parte Skinner*, 2 USPQ2d 1788, 1788-89 (B.P.A.I.

1986). The reference(s) *must teach each and every aspect of the claimed invention* either explicitly or impliedly, and the burden is on the Office to present a *prima facie* case of unpatentability.

The Office relies on Amiri et al. for teaching most limitations, and admits that "Amiri fails to disclose that the protected data is included in the release and grant messages and that the protected data is copied from the release to the grant message." Final Office action, page 11, lines 9-10. Appellants agree. The Office then relies on Yun et al. for these teachings.

First, Appellants note that claim 6 requires that the operations be performed by the lock manager. As discussed *supra* on pages 12 and 13, the manager of the locks does not receive, nor have access to the data protected by the manager of the locks; and therefore, neither teaches nor suggests performs the recited limitations for which the Office relies upon Yun et al. Yun et al.'s process having the lock and processing the data is not a lock manager. For at least these reasons, the Office has failed to present a *prima facie* rejection.

Next, the Office action never presents a rejection for the release containing the protected data. Rather, it merely presents a rejection for the grant message. For at least this reason, the Office has failed to present a *prima facie* rejection.

Next, a "release" is a known entity in the art. A release is sent from a process to the lock manager to inform the lock manager that it is done with the protected data, so that the lock manager can grant access to the protected data to a currently requesting process. There is no teaching in Yun et al. that when a process forwards the diffs to the acquiring process, it sends a release including the protected data to the lock manager. Rather, Yun et al., as further explained by Y. Zhou et al, teaches that the manager of the lock sends the grant for the acquiring process to the process holding the lock. The lock holder, when completed with data, then typically creates a message with the diffs and forwards to the acquiring process (which was determined by the lock manager). Yun et al. neither teaches nor suggest a release message with the protected data being sent to the lock manager. For at least this reason, the Office has failed to present a *prima facie* rejection.

Moreover, claim 6 recites that the next requester to be granted the lock is identified in response to said receiving the release of the lock. As discussed, Yun et al. teaches a different protocol. Again, the lock manager identifies and grants a next requester access to the data while it is still locked by the current process, and sends this grant to the current process. When the current process completes, it sends the diffs as a lock grant message to the granted process (determined by the lock manager). This neither teaches nor suggests the recited limitation that the next requester to be granted the lock is identified in response to said receiving the release of the lock. For at least these reasons, the Office has failed to present a *prima facie* rejection.

Finally, Yun et al. neither teaches nor suggests copying the protected data from the release into a grant message. Again, Yun et al. neither teaches nor suggests a release with the protected data sent to the lock manager. Next, the Office's rejection of claim 6 includes that "in order for the protected data/diffs to move from the release to the grant message, it must be copied there." Appellants make a "demand for evidence." This statement in the Office action is not a teaching of Yun et al. Rather, Yun et al. teaches that the releasing process creates the diff message from data it has processed (e.g., read, written) to the protected memory location. There is no release message sent to the locking manager with the protected data, and there is no copying by the locking manager (or even a process P0-P4) of the protected data from a release to the grant. For at least these reasons, the Office has failed to present a *prima facie* rejection.

Appellants further traverse the combination of Yun et al. with Amiri et al. that would render a claim unpatentable as being obvious as Yun et al. communicating of data is from a releasing process to an acquiring process; and Yun et al. or Amiri et al., alone or in combination, neither teach nor suggest protected data being communicated or even touched by a lock manager as recited in every pending claim.

For at least these reasons, Appellants request the Board reverse the rejections of claims 6-11.

(vii)(c) The Office fails to present a *prima facie* rejection of any of claims 12-26, as the prior art of record, alone or in combination, neither teaches nor suggests a lock manager sending a first grant message to the first requester, the first grant message not including said protected data, and in response to identifying one or more requesters is waiting for the lock after the first requester, including an indication to return said protected data in the first grant message; and the lock manager receiving a first release message corresponding to the first grant message for the lock from the first requester, the first release message including said protected data.

Group: claims 12-26.

Representative claim: independent claim 12.

Independent claim 12 recites:

"12. A method performed by a lock manager controlling access to protected data maintained in native storage independent of the lock manager, wherein the lock manager does not access said protected data from said native storage, the method comprising:

receiving locking requests for a lock controlling access to said protected data from a first requester and a second requester;

sending a first grant message to the first requester, the first grant message not including said protected data, and in response to identifying one or more requesters is waiting for the lock after the first requester, including an indication to return said protected data in the first grant message; and

receiving a first release message corresponding to the first grant message for the lock from the first requester, the first release message including said protected data."

Claims 12-26 stand rejected per the final Office mailed July 23, 2007, as being obvious under 35 USC § 103(a) over the Amiri et al., "Highly Concurrent Shared Storage" article in view of the Yun et al., "An Efficient Lock Protocol for Home-based Lazy Release Consistency" article. Appellants respectfully traverse the rejection of these claims, as the prior

art of record neither teaches nor suggests all recited limitations of any of these claims. Note, independent claim 17 and its dependent claims 18-21 and independent claim 22 and its dependent claims 23-26 are different claim formats directly based on independent claim 12 and its dependent claims 13-16, respectively.

It is well-established law that the burden is on the Office to initially present a *prima facie* unpatentability rejection, before Applicant has any burden of proof of disproving any application of a cited reference against a claim. *In re Warner*, 379 F2d. 1011, 1016, 154 USPA 173, 177 (C.C.P.A. 1967); *Ex parte Skinner*, 2 USPQ2d 1788, 1788-89 (B.P.A.I. 1986). The reference(s) *must teach each and every aspect of the claimed invention* either explicitly or impliedly, and the burden is on the Office to present a *prima facie* case of unpatentability.

The Office relies on Amiri et al. for teaching most limitations, and admits that "Amiri fails to disclose in response to identifying one or more requesters is waiting for the lock after the first requester, including an indication to return said protected data in the first grant message and the first release message including said protected data." Final Office action, page 14, second paragraph. Appellants agree. The Office then relies on Yun et al. for these teachings.

First, Appellants note that claim 12 requires that *all of the recited operations* be performed by the lock manager. As discussed *supra* on pages 12 and 13, the manager of the locks does not receive, nor have access to the data protected by the manager of the locks; and therefore, neither teaches nor suggests performs the recited limitations for which the Office relies upon Yun et al. Yun et al.'s process having the lock and processing the data is not a lock manager. Additionally, a rejection relying on a process (e.g., Yun et al.'s P0, P1, P2, P3) that manipulates the data conflicts with the claim limitation that the lock manager (performing all of these operations) does not access said protected data from said native storage.

Claim 12 recites that the lock manager sends the first grant message including an indication to return (i.e., to the lock manger) the protected data in response to identifying one

or more requesters is waiting for the lock after the first requester. The Office's rejection is *non sequitur* to all of these limitations. The statement of rejection states that the releasing process sends *the actual protected data*; while the claim limitation requires the sending of *an indication to return the protected data*.

Next, a "release" is a known entity in the art. A release is sent from a process to the lock manager to inform the lock manager that it is done with the protected data, so that the lock manager can grant access to the protected data to a currently requesting process. There is no teaching in Yun et al. that when a process forwards the diffs to the next acquiring process, it sends a release including the protected data to the lock manager. Rather, Yun et al., as further explained by Y. Zhou et al, teaches that the manager of the lock sends the grant for the acquiring process to the process holding the lock. The lock holder, when completed with data, then typically creates a message with the diffs and forwards to the acquiring process. Yun et al. neither teaches nor suggest a release message with the protected data being sent to the lock manager.

Not only does Yun et al. operate differently as discussed *supra*, the statement of the rejection in the Office action fails to address the structure required by the claim. The claim limitations require the same entity (the lock manager) *that sends the grant* to the first requester *to receive the first release corresponding to the first grant* from the first requester, with the first release including the protected data. The rejection in the Office action equates the sending of the diffs by a releasing process to an acquiring process as the sending of the grant message. Therefore, for a coherent rejection, the acquiring process would need to send a release to the releasing process. This rejection therefore makes no sense.

Appellants further traverse the combination of Yun et al. with Amiri et al. that would render a claim unpatentable as being obvious as Yun et al. communicating of data is from a releasing process to an acquiring process; and Yun et al. or Amiri et al., alone or in combination, neither teach nor suggest protected data being communicated or even touched by a lock manager as recited in every pending claim.

For at least this reason, the Office has failed to present a *prima facie* rejection of independent claim 12.

For at least these reasons, Appellants request the Board reverse the rejections of claims 12-26.

(vii)(d) The Office fails to present a prima facie rejection of any of claims 4, 8, 9, 11, 13 16, 18-21, and 23-26.

Group: claims 4, 8, 9, 11, 13 16, 18-21, and 23-26.

Representative claim: independent claim 4.

Claim 4 recites "[t]he apparatus of claim 1, wherein each of said communicated grants includes an indication of whether or not said protected data is requested to be sent to the lock manager with a corresponding release of the lock."

The Office action relies on Yun et al. for this teaching, by apparently making an inherency argument based on speculation of the operation of Yun et al., reasoning that a grant message to a process must include such an indication so it can determine whether send the data to a next acquiring process or write it back to storage. Final Office action at page 8.

"In relying upon the theory of inherency, the examiner must provide a basis in fact and/or technical reasoning to reasonably support the determination that the allegedly inherent characteristic necessarily flows from the teachings of the applied prior art." *Ex parte Levy*, 17 USPQ2d 1461, 1464 (Bd. Pat. App. & Inter. 1990) (emphasis in original) . Inherent means it *must* occur. The fact that a certain result or characteristic *may* occur or be present in the prior art is not sufficient to establish the inherency of that result or characteristic.

The statement of the rejection in the Office action includes "... The preceding text along with Figure 2 clearly indicates that if no other processes are requesting the lock, that the protected data is written back to storage, rather than being forwarded to a next acquiring process. In order to make this determination and perform this operation, an indication of whether or not to forward the protected data must be including in the grant message...."

Appellants interpret that the grant message to which the rejection refers is the one granting the current process access to the data (i.e., giving it the lock), as if there is no waiting process (which is the premise of the Office's argument), there will not be a subsequent grant received by the process.

There are alternatives to the mechanism stated in the Office action, including most likely how Yun et al. actually operates: if there is a process waiting for the lock, the lock manager sends a grant message for the waiting process to the process holding the lock; else it does not send a message. *See*, the excerpt from Y. Zhou et al, *supra*. Therefore, a process holding the lock knows whether or not there is a waiting process based on whether or not it receives a grant message while it has the lock for the data. Hence, the Office's argument that "an indication of whether or not to forward the protected data must be including in the grant message" is wrong.

Furthermore, the explanation in the Office action does not make sense as it would not operate as apparently intended by Yun et al. when there is no other outstanding request for the data when process acquires the lock to the data (i.e., when it receives the grant message granting it access to the lock), and a request is made subsequently for the data by another process while the original process still has the lock. In this scenario, Yun et al. as interpreted by the Office would not know of the subsequent request for the data, as the indication of whether or not to forward the protected data must be included in the grant message giving the lock to the current process.

For at least these additional reasons, the Office fails to present a *prima facie* rejection of dependent claim 4 as the Office relies on Yun et al., which neither teaches nor suggests this recited limitation.

Additionally, as discussed *supra*, the Yun et al. neither teaches nor suggests that the manager of the lock ever actually receives the protected data; and therefore, there is no need for, nor is an indication of whether or not to send the protected to the lock manager taught by Yun et al., as the manager of the lock

For at least these reasons, Appellants request the Board reverse the rejections of dependent claims 4, 8, 9, 11, 13 16, 18-21, and 23-26.

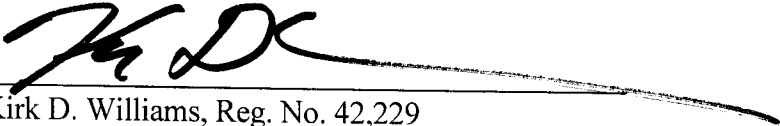
In re WILLIAMS JR. ET AL., Application No. 10/811,044
APPEAL BRIEF

FINAL REMARKS. In view of the above remarks, for at least the reasons presented herein, and that the prior art of record neither teaches nor suggests all the elements/limitations of any pending claim, all pending claims are believed to be allowable over the prior art of record, the application is considered in good and proper form for allowance. Appellants respectfully request all claim rejections be reversed and all claims be allowed. Additionally, Appellants request the Office withdraw all rejections and/or objections and allow the case in response to this reply to the final Office action.

Date: April 21, 2008

Respectfully submitted,
The Law Office of Kirk D. Williams

By



Kirk D. Williams, Reg. No. 42,229
One of the Attorneys for Appellant
CUSTOMER NUMBER 26327
The Law Office of Kirk D. Williams
PO Box 80238-8538, Denver, CO 80206-8538
303-282-0151 (telephone), 303-778-0748 (facsimile)

(viii) CLAIMS APPENDIX

1. An apparatus for protecting data using locks, the apparatus comprising:
a lock manager configured to control access via a lock to protected data maintained in native storage independent of the lock manager, wherein the lock manager does not access said protected data from said native storage; and
a plurality of requesters;
wherein the lock manager is configured to receive lock requests for the lock from each of the plurality of requesters, and to selectively grant said lock requests which includes communicating grants from the lock manager to the plurality of requesters, wherein at least one of said communicated grants includes said protected data.
2. The apparatus of claim 1, wherein at least one of said communicated grants does not include said protected data.
3. The apparatus of claim 1, wherein each of said communicated grants includes an indication of whether or not said protected data is being communicated therewith.
4. The apparatus of claim 1, wherein each of said communicated grants includes an indication of whether or not said protected data is requested to be sent to the lock manager with a corresponding release of the lock.
5. The apparatus of claim 1, wherein each of said lock requests includes an indication of whether or not the corresponding one of the plurality of requesters will accept said protected data from the lock manager.

6. A method performed by a lock manager controlling access to protected data maintained in native storage independent of the lock manager, wherein the lock manager does not access said protected data from said native storage, the method comprising:

receiving a release of a lock for use in controlling access to said protected data, the release including said protected data;

identifying a next requester to be granted the lock in response to said receiving the release of the lock;

copying said protected data from the release into a grant message; and

sending the grant message to the next requester, the grant message including said protected data.

7. The method of claim 6, wherein the grant message includes an indication of that said protected data is requested to be sent to the lock manager in a release message corresponding to the grant message if another requester is waiting for the lock, else an indication that said protected data is not requested to be sent to the lock manager in the release message.

8. A computer-readable medium tangibly storing thereon computer-executable instructions for performing steps by a lock manager for controlling access to protected data maintained in native storage independent of the lock manager, wherein the lock manager does not access said protected data from said native storage, said steps comprising:

receiving a release of a lock for use in controlling access to said protected data, the release including said protected data;

identifying a next requester to be granted the lock in response to said receiving the release of the lock;

copying said protected data from the release into a grant message; and

sending the grant message to the next requester, the grant message including said protected data.

9. The computer-readable medium of claim 8, wherein the grant message includes an indication of that said protected data is requested to be sent to the lock manager in a release message corresponding to the grant message if another requester is waiting for the lock, else an indication that said protected data is not requested to be sent to the lock manager in the release message.

10. A lock manager controlling access to protected data maintained in native storage independent of the lock manager, wherein the lock manager does not access said protected data from said native storage, the lock manager comprising:

means for receiving a release of a lock for use in controlling access to said protected data, the release including said protected data;

means for identifying a next requester to be granted the lock in response to said receiving the release of the lock;

means for copying said protected data from the release into a grant message and for sending the grant message to the next requester.

11. The lock manager of claim 10, means for including in the grant message an indication of that said protected data is requested to be sent to the lock manager in a release message corresponding to the grant message if another requester is waiting for the lock, else an indication that said protected data is not requested to be sent to the lock manager in the release message.

12. A method performed by a lock manager controlling access to protected data maintained in native storage independent of the lock manager, wherein the lock manager does not access said protected data from said native storage, the method comprising:

receiving locking requests for a lock controlling access to said protected data from a first requester and a second requester;

sending a first grant message to the first requester, the first grant message not including said protected data, and in response to identifying one or more requesters is waiting for the lock after the first requester, including an indication to return said protected data in the first grant message; and

receiving a first release message corresponding to the first grant message for the lock from the first requester, the first release message including said protected data.

13. The method of claim 12, comprising sending a second grant message to the second requester, the second grant message including said protected data, and an indication of whether or not to send said protected data in a second release message.

14. The method of claim 13, wherein the second grant message includes an indication to send said protected data in the second release message in response to identifying another requestor is waiting for access to the lock.

15. The method of claim 13, wherein the second grant message includes an indication not to send said protected data in the second release message in response to identifying another requestor is not waiting for access to the lock.

16. The method of claim 13, wherein the second grant message includes an indication not to send said protected data in the second release message; and the method comprises in response to said indication not to send said protected data in the second release message, the second requester storing said protected data and not including said protected data in the second release message.

17. A computer-readable medium tangibly storing thereon computer-executable instructions for performing steps by a lock manager for controlling access to protected data maintained in native storage independent of the lock manager, wherein the lock manager does not access said protected data from said native storage, said steps comprising:

receiving locking requests for a lock controlling access to said protected data from a first requester and a second requester;

sending a first grant message to the first requester, the first grant message not including said protected data, and in response to identifying one or more requesters is waiting for the lock after the first requester, including an indication to return said protected data in the first grant message; and

receiving a first release message corresponding to the first grant message for the lock from the first requester, the first release message including said protected data.

18. The computer-readable medium of claim 17, wherein said steps comprise sending a second grant message to the second requester, the second grant message including said protected data, and an indication of whether or not to send said protected data in a second release message.

19. The computer-readable medium of claim 18, wherein the second grant message includes an indication to send said protected data in the second release message in response to identifying another requestor is waiting for access to the lock.

20. The computer-readable medium of claim 18, wherein the second grant message includes an indication not to send said protected data in the second release message in response to identifying another requestor is not waiting for access to the lock.

21. The computer-readable medium of claim 18, wherein the second grant message includes an indication not to send said protected data in the second release message; and said steps comprise in response to said indication not to send said protected data in the second release message, the second requester storing said protected data and not including said protected data in the second release message.

22. A lock manager controlling access to protected data maintained in native storage independent of the lock manager, wherein the lock manager does not access said protected data from said native storage, the lock manager comprising:

means for receiving locking requests for a lock controlling access to said protected data from a first requester and a second requester;

means for sending a first grant message to the first requester, the first grant message not including said protected data, and in response to identifying one or more requesters is waiting for the lock after the first requester, including an indication to return said protected data in the first grant message; and

means for receiving a first release message for the lock from the first requester, the first release message including said protected data.

23. The lock manager of claim 22, comprising means for sending a second grant message to the second requester, the second grant message including said protected data, and an indication of whether or not to send said protected data in a second release message.

24. The lock manager of claim 23, comprising means for including in the second grant message an indication to send said protected data in the second release message in response to identifying another requestor is waiting for access to the lock.

25. The lock manager of claim 23, comprising means for including in the second grant message an indication not to send said protected data in the second release message in response to identifying another requestor is not waiting for access to the lock.

26. The lock manager of claim 23, comprising: means for including in the second grant message an indication not to send said protected data in the second release message; and means for the second requester to store said protected data and not to include said protected data in the second release message in response to said indication not to send said protected data in the second release message.

(ix) EVIDENCE APPENDIX

NONE.

In re WILLIAMS JR. ET AL., Application No. 10/811,044
APPEAL BRIEF

(x) RELATED PROCEEDINGS APPENDIX

NONE.